

# FAST ALGORITHMS FOR BAYESIAN UNCERTAINTY QUANTIFICATION IN LARGE-SCALE LINEAR INVERSE PROBLEMS BASED ON LOW-RANK PARTIAL HESSIAN APPROXIMATIONS\*

V. AKÇELIK<sup>†</sup>, H. P. FLATH<sup>‡</sup>, O. GHATTAS<sup>§</sup>, J. HILL<sup>¶</sup>, B. VAN BLOEMEN WAANDERS<sup>||</sup>  
, AND L. C. WILCOX<sup>\*\*</sup>

**Abstract.** We consider the problem of estimating the uncertainty in large-scale linear statistical inverse problems with high-dimensional parameter spaces within the framework of Bayesian inference. When the noise and prior probability densities are Gaussian, the solution to the inverse problem is also Gaussian, and is thus characterized by the mean and covariance matrix of the posterior probability density. Unfortunately, explicitly computing the posterior covariance matrix requires as many forward solutions as there are parameters, and is thus prohibitive when the forward problem is expensive and the parameter dimension is large. However, for many ill-posed inverse problems, the Hessian matrix of the data misfit term has a spectrum that collapses rapidly to zero. We present a fast method for computation of an approximation to the posterior covariance that exploits the low-rank structure of the preconditioned (by the prior covariance) Hessian of the data misfit. Analysis of an infinite-dimensional model convection-diffusion problem, and numerical experiments on large-scale 3D convection-diffusion inverse problems with up to 1.5 million parameters, demonstrate that the number of forward PDE solves required for an accurate low-rank approximation is independent of the problem dimension. This permits scalable estimation of the uncertainty in large-scale ill-posed linear inverse problems at a small multiple (independent of the problem dimension) of the cost of solving the forward problem.

**Key words.** large-scale statistical inverse problem, Bayesian inference, uncertainty quantification, fast algorithms, low-rank approximation, Lanczos, Hessian, convection-diffusion, contaminant transport

**AMS subject classifications.** 35Q62 35Q93 35Q99 76R99 62F15 65C60 65M32

**1. Introduction.** Uncertainty quantification for inverse problems governed by partial differential equations (PDEs) with high-dimensional parameter spaces poses considerable computational difficulties and is often intractable. We consider the case

---

\*This research partially supported by NSF grants OPP-0941678, DMS-0724746, CNS-0619838, CNS-0540372, and CCF-0427985; DOE grants DE-SC0002710, DE-FG02-08ER25860, DE-FC52-08NA28615, and DE-FC02-06ER25782; and AFOSR grants FA9550-09-1-0608 and FA9550-07-1-0480. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company for the United States Department of Energys National Nuclear Security Administration under contract DE-AC04-94AL85000. The second author's research was partially performed under an appointment to the Department of Homeland Security (DHS) Scholarship and Fellowship Program, administered by the Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy (DOE) and DHS. ORISE is managed by Oak Ridge Associated Universities (ORAU) under DOE contract number DE-AC05-06OR23100. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DHS, DOE, or ORAU/ORISE.

<sup>†</sup>Stanford Linear Accelerator Center (now at Exxon-Mobil Research and Engineering Company)

<sup>‡</sup>Institute for Computational Engineering & Sciences and Computational Science, Engineering, & Mathematics graduate program, The University of Texas at Austin. Author to whom correspondence should be sent at [pflath@ices.utexas.edu](mailto:pflath@ices.utexas.edu)

<sup>§</sup>Institute for Computational Engineering & Sciences, Jackson School of Geosciences, and Department of Mechanical Engineering, The University of Texas at Austin

<sup>¶</sup>Uncertainty and Optimization Department, Sandia National Laboratories (now in Computational Mathematics Group, Oak Ridge National Laboratory)

<sup>||</sup>Uncertainty and Optimization Department, Sandia National Laboratories (now in Applied Mathematics and Applications Department, Sandia National Laboratories)

<sup>\*\*</sup>Institute for Computational Engineering & Sciences, The University of Texas at Austin

of a linear ill-posed inverse problem within the framework of Bayesian inference [19,23]. The relevant background is summarized in §2. For this class of problem, an additive Gaussian noise model and a Gaussian prior density on the model parameters will result in a Gaussian posterior density for the model parameters. Thus, the statistical solution of the inverse problem, i.e. the posterior density, can be completely described by its mean and covariance matrix. The mean is given by maximizing the posterior density, which leads to a linear least squares optimization problem whose structure is identical to that of a properly weighted deterministic inverse problem. Thus, state-of-the-art numerical algorithms from large-scale deterministic inverse problems can be exploited to find the mean. On the other hand, the posterior covariance matrix  $\mathbf{\Gamma}_{\text{post}}$  is given by the inverse of the Hessian matrix of this least squares function. Due to the high-dimensionality of the parameter space and expense of solving the underlying forward and adjoint PDEs needed to construct the Hessian, direct computation of  $\mathbf{\Gamma}_{\text{post}}$  is usually intractable. Moreover, the explicit inverse of the Hessian is not required for deterministic inversion, and thus one cannot appeal to deterministic techniques for directly computing  $\mathbf{\Gamma}_{\text{post}}$ .

The goal of this paper is to present and analyze a fast method for accurate estimation of the posterior covariance matrix  $\mathbf{\Gamma}_{\text{post}}$  by exploiting the structure of the Hessian. Under the conditions described above, the Hessian has two components, one stemming from the weighted data misfit, and one from the prior information on the parameters. The posterior covariance can then be written as  $\mathbf{\Gamma}_{\text{post}} = \left( \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} + \mathbf{\Gamma}_{\text{prior}}^{-1} \right)^{-1}$ , where  $\mathbf{\Gamma}_{\text{prior}}$  is the prior covariance matrix,  $\mathbf{\Gamma}_{\text{noise}}$  is the noise covariance matrix, and  $\mathbf{A}$  is the parameter-to-observable map, which involves the inverse of the forward operator, and is therefore usually dense. As will be illustrated in §5, explicit computation of  $\mathbf{A}$  requires as many forward PDE solves as there are parameters. Thus, for high-dimensional parameter spaces and expensive-to-solve PDEs, computing  $\mathbf{\Gamma}_{\text{post}}$  by construction and inversion of the Hessian is prohibitive. However, for many ill-posed inverse problems, the spectrum of  $\mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A}$  (the data misfit portion of the Hessian) decays rapidly, reflecting the fact that the data provide information on only a low-dimensional subspace of the parameter field. One would like to employ a low-rank approximation of this operator (e.g., via a Lanczos method) to overcome the prohibitive nature of direct computation of  $\mathbf{\Gamma}_{\text{post}}$ . However, the retained modes should be informed by not only the data, but also the prior information, since modes that can be inferred from the data might be nullified by the influence of the prior. Therefore, we rearrange the expression for  $\mathbf{\Gamma}_{\text{post}}$  to include the effects of the prior in the data misfit term, thereby filtering the data through the prior before invoking the low-rank approximation.

We argue in subsequent sections that for many ill-posed inverse problems, an accurate low-rank spectral approximation of the “prior-preconditioned Hessian of the data misfit,”  $\mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2}$  can be computed at a cost that is a small multiple (independent of the parameter dimension) of the cost of solving the underlying forward PDEs. Once the low-rank approximation has been made, we can compute desired quantities from the approximate posterior at a cost that depends only linearly on problem dimension. These features result in a scalable method for estimating uncertainty in large-scale linear statistical inverse problems.

In the remainder of this section, we review pertinent literature and describe the differences with our approach and the contributions of this paper. Much of the prior work for low-rank approximations in inverse problems has occurred in the context of seismic inverse problems, where the forward operators can be expensive, and the

heterogeneity of the earth requires a high-dimensional parametrization. When the forward problem is sufficiently inexpensive to solve, one can explicitly construct the Hessian matrix (at a cost of  $n$  forward problem solutions, where  $n$  is the number of model parameters) and invert it to obtain the posterior covariance matrix. For example, as early as 1993, [25] were able to explicitly compute the posterior covariance in the Bayesian framework for a global seismic tomography problem using ray tracing as the forward model for up to 12,496 model parameters. Similarly, when the parametrization is low-dimensional, one can tractably construct and invert the Hessian matrix to find the posterior covariance. For example, [15] and [16] obtain Bayesian solutions to flat-layered earth seismic inverse problems with up to 300 parameters. Such explicit construction of the Hessian does not scale as the number of parameters, or the complexity of the forward problem, increase. On the other hand, a number of references invoke low-rank approximations to make inversion tractable for large-scale problems. Examples include [8, 9, 11, 14, 20, 21, 22, 24, 27, 28, 29], in which the low rank approximations are effected through truncated iterative methods such as Lanczos and block Lanczos, conjugate gradients, and LSQR. In all of these cases, however, the Bayesian framework is not employed, and instead the goal is to determine the so-called resolution matrix  $\mathbf{A}^\dagger \mathbf{A}$  [4], which is a deterministic attempt to quantify resolving power of the inversion.

Our approach differs from previous work in the following ways:

- We adopt the full Bayesian statistical framework to estimate uncertainties in solution of inverse problems, thus permitting a broader interpretation of the regularization as a model prior; importantly, our low-rank approximation addresses both data misfit and model prior terms, which leads to a systematic method to choose how to truncate the low-rank approximation (see §3).
- Our method address very large-scale inverse problems for which both the forward problem and the model parametrization are large-scale. For example, we apply the method to a 3D implicitly-solved convection-diffusion atmospheric transport inverse problem posed on a complex geometry domain with finite element-based discretization of the initial condition with up to 1.5 million inversion parameters (see §5).
- For the 3D inverse convection-diffusion problem, we provide a systematic investigation of the effects of truncation of the low-rank estimate on the accuracy of the posterior, and we assess the scalability of the method with increasing parameter dimension, demonstrating that the number of forward PDE solves required by an accurate low-rank-based approximation is independent of the problem dimension (see §5). We also provide an analysis of the spectrum of the Hessian of a 1D model convection-diffusion inverse problem, the results of which support the numerical experiments (see §4).

We begin with an introduction to the Bayesian framework in §2. Next, in §3, we describe the low-rank algorithm for approximating the posterior covariance matrix. §4 introduces the specific inverse problem that is the subject of our numerical experiments: inversion for the initial concentration field of an atmospheric contaminant governed by a 3D convection-diffusion PDE in a model city geometry, based on sensor observations. This sections also assesses the efficacy of the low-rank approach through spectral analysis of a 1D model problem. We study the performance of this algorithm in §5 for large-scale instances of the 3D convection-diffusion inverse problem. We investigate the structure of the spectrum of the prior-preconditioned Hessian of the data misfit in §5.1 and its dependence on physical and experimental parameters in

§5.2. In §5.3–5.4, we provide interpretation and error analysis of the approximated variance. Finally, §5.5 demonstrates the scalability of the low-rank algorithm as a function of the number of parameters.

**2. Bayesian framework for statistical inverse problems.** In this section, we present the Bayesian framework for statistical inverse problems beginning with the general case of Bayes’ theorem and continuing with the special case of a linear forward model and Gaussian noise and prior uncertainties. The resulting posterior probability density will be Gaussian, with mean given by the solution of a particular least squares inverse problem, and the covariance given by the inverse of the Hessian matrix of the least squares function. Our focus here is on inverse problems that are governed by large-scale forward models, as result from appropriate discretizations of PDEs. Unfortunately, for large-scale statistical inverse problems (as exemplified by the inverse convection-diffusion problem we target here), a straightforward computation of this covariance matrix is entirely prohibitive. Hence, in §3, we present a fast method for estimating the covariance based on a low rank spectral estimate that approximates the structure of the Hessian.

**2.1. General formulation.** Non-uniqueness is a central feature of ill-posed inverse problems: multiple values of the parameters may be consistent with the observations. The least-squares minimization approach to ill-posed inverse problems invokes a so-called regularization terms to effectively select among the multiple parameter values the one that has largest regularity (in an appropriate norm), resulting in a single—deterministic—estimate of the unknown parameters (see, e.g., [26]). A Bayesian estimate of the unknown, on the other hand, is a probability density that suggests the credibility of any given point estimate (see, e.g., [19,23]). In the Bayesian approach, we view all parameters as random variables and write the parameter-to-observable map  $g : \mathbb{R}^n \times \mathbb{R}^k \rightarrow \mathbb{R}^m$  as

$$Y = g(X, E),$$

where  $X$ ,  $Y$ , and  $E$  are random variables. The variable  $x \in \mathbb{R}^n$  is a realization of the random variable  $X$  and the vector of model parameters to be recovered,  $e \in \mathbb{R}^k$  is a realization of the random variable  $E$  and the vector of errors (due to both model errors and observation noise), and  $y \in \mathbb{R}^m$  is the realization of random variable  $Y$  and vector of observables with  $\mathbf{y}_{\text{obs}}$  the actual observation values. We choose the following probability density functions (pdf’s): the probability density  $\pi_{\text{noise}} : \mathbb{R}^k \rightarrow \mathbb{R}$ , which describes the modeling error and observation noise; the prior probability density  $\pi_{\text{prior}} : \mathbb{R}^n \rightarrow \mathbb{R}$ , which describes additional information about the parameters  $X$ ; and the likelihood function  $\pi(\mathbf{y}|\mathbf{x})$ , which describes the relationship between the observables  $\mathbf{y}$  and the unknown model parameters  $\mathbf{x}$ .

From Bayes’ theorem, the prior probability density, the likelihood function, and the data can be combined to form the posterior probability density  $\pi_{\text{post}} : \mathbb{R}^n \rightarrow \mathbb{R}$  on the model parameters  $X$ :

$$\begin{aligned} \pi_{\text{post}}(\mathbf{x}) &:= \pi(\mathbf{x}|\mathbf{y}_{\text{obs}}) \\ &= \frac{\pi_{\text{prior}}(\mathbf{x}) \pi(\mathbf{y}_{\text{obs}}|\mathbf{x})}{\pi(\mathbf{y}_{\text{obs}})} \\ &\propto \pi_{\text{prior}}(\mathbf{x}) \pi(\mathbf{y}_{\text{obs}}|\mathbf{x}). \end{aligned}$$

To generate the likelihood function, we use the pdf  $\pi_{\text{noise}}(\mathbf{e})$ . Here we will assume

additive noise, so that the parameter-to-observable map is

$$Y = f(X) + E,$$

where  $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$  and the noise  $E \in \mathbb{R}^m$  reflects both the modeling error of  $f$  and observation noise. Thus  $E = Y - f(X)$ . We assume that  $X$  and  $E$  are statistically independent. Therefore,

$$\pi_{\text{noise}}(\mathbf{e}) = \pi_{\text{noise}}(\mathbf{y}_{\text{obs}} - f(\mathbf{x})),$$

so that Bayes' theorem states

$$\pi_{\text{post}}(\mathbf{x}) \propto \pi_{\text{prior}}(\mathbf{x}) \pi_{\text{noise}}(\mathbf{y}_{\text{obs}} - f(\mathbf{x})).$$

**2.2. Gaussian linear case.** In the case that the prior probability density of  $X$  and the probability density of error  $E$  are both Gaussian, further simplifications can be made. The prior and noise pdf's can be written in the form

$$\begin{aligned} \pi_{\text{prior}}(\mathbf{x}) &\propto \exp\left(-\frac{1}{2}(\mathbf{x} - \bar{\mathbf{x}}_{\text{prior}})^T \mathbf{\Gamma}_{\text{prior}}^{-1}(\mathbf{x} - \bar{\mathbf{x}}_{\text{prior}})\right), \\ \pi_{\text{noise}}(\mathbf{e}) &\propto \exp\left(-\frac{1}{2}(\mathbf{e} - \bar{\mathbf{e}})^T \mathbf{\Gamma}_{\text{noise}}^{-1}(\mathbf{e} - \bar{\mathbf{e}})\right), \end{aligned}$$

where  $\bar{\mathbf{x}}_{\text{prior}} \in \mathbb{R}^n$  is the mean of the model parameter prior pdf,  $\bar{\mathbf{e}} \in \mathbb{R}^m$  is the mean of the noise pdf,  $\mathbf{\Gamma}_{\text{prior}} \in \mathbb{R}^{n \times n}$  is the covariance matrix of the prior pdf, and  $\mathbf{\Gamma}_{\text{noise}} \in \mathbb{R}^{m \times m}$  is the covariance matrix of the noise pdf. Restating Bayes' theorem with these Gaussian uncertainties, we find that

$$\pi_{\text{post}}(\mathbf{x}) \propto \exp\left(-\frac{1}{2}\|\mathbf{x} - \bar{\mathbf{x}}_{\text{prior}}\|_{\mathbf{\Gamma}_{\text{prior}}^{-1}}^2 - \frac{1}{2}\|\mathbf{y}_{\text{obs}} - f(\mathbf{x}) - \bar{\mathbf{e}}\|_{\mathbf{\Gamma}_{\text{noise}}^{-1}}^2\right).$$

As is clear from this expression, despite the choice of prior and noise probability densities as Gaussian, the posterior probability density need not be Gaussian, due to the nonlinearity of  $f(\mathbf{x})$ .

Let us now assume that the parameter-to-observable map is linear, i.e.

$$f(X) = \mathbf{A}X.$$

Here,  $\mathbf{A} \in \mathbb{R}^{m \times n}$  is the linear operator that maps parameters  $\mathbf{x}$  to observables  $\mathbf{y}$ , via the solution of a large-scale discretized PDE problem. In this case,  $\pi_{\text{post}}(\mathbf{x})$  is also Gaussian, with mean  $\bar{\mathbf{x}}_{\text{post}} \in \mathbb{R}^n$  given by the maximum *a posteriori* (MAP) point, i.e.,

$$\bar{\mathbf{x}}_{\text{post}} = \arg \max_{\mathbf{x}} \pi_{\text{post}}(\mathbf{x}).$$

Finding the MAP point is equivalent to solving a weighted least squares optimization problem, i.e.,

$$\bar{\mathbf{x}}_{\text{post}} = \arg \min_{\mathbf{x}} \left( \frac{1}{2}\|\mathbf{y}_{\text{obs}} - \mathbf{A}\mathbf{x} - \bar{\mathbf{e}}\|_{\mathbf{\Gamma}_{\text{noise}}^{-1}}^2 + \frac{1}{2}\|\mathbf{x} - \bar{\mathbf{x}}_{\text{prior}}\|_{\mathbf{\Gamma}_{\text{prior}}^{-1}}^2 \right), \quad (1)$$

which amounts to solving a regularized deterministic inverse problem, where  $\mathbf{\Gamma}_{\text{prior}}^{-1}$  plays the role of the regularization operator, and  $\mathbf{\Gamma}_{\text{noise}}^{-1}$  is a weighting of the data misfit

term. The covariance matrix of the posterior pdf of model parameters,  $\mathbf{\Gamma}_{\text{post}} \in \mathbb{R}^{n \times n}$ , is given simply by the inverse of the Hessian matrix of the least squares objective function, i.e.,

$$\mathbf{\Gamma}_{\text{post}} = \left( \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} + \mathbf{\Gamma}_{\text{prior}}^{-1} \right)^{-1}. \quad (2)$$

So, in summary,

$$\pi_{\text{post}}(\mathbf{x}) = \mathcal{N}(\bar{\mathbf{x}}_{\text{post}}, \mathbf{\Gamma}_{\text{post}}).$$

**3. Low-rank approximation.** While the mean of the posterior density of the parameters can be found readily by solving (1) using state-of-the-art techniques from large-scale PDE-constrained optimization (see e.g. the review in [3] or the edited volume [10]), computing the posterior covariance matrix of the parameters presents a challenge. Although it is easy to write an expression such as (2), explicit construction of the Hessian (let alone its inverse), which is dense and of order of the number of parameters, would require at least as many forward solves as the number of parameters to compute the parameter-to-observable map  $\mathbf{A}$ . This is prohibitive for high-dimensional parameter spaces, as arise when an infinite-dimensional parameter field is discretized. Moreover, for large-scale problems we often cannot store the entire posterior covariance matrix, and instead must resort to extracting partial information (such as its action on a vector or the variance).

Instead, we take advantage of the structure of the Hessian. We begin by recognizing that, for many ill-posed inverse problems, the Hessian of the data misfit,  $\mathbf{H}_{\text{misfit}} \in \mathbb{R}^{n \times n}$ , where

$$\mathbf{H}_{\text{misfit}} \stackrel{\text{def}}{=} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A},$$

behaves like the discretization of a compact operator. The range space thus is effectively finite-dimensional, and the eigenvalues decay, often rapidly, to zero. We can exploit this structure to construct fast algorithms for approximating the inverse of the Hessian. Rearranging the expression for  $\mathbf{\Gamma}_{\text{post}}$  in (2) to factor out  $\mathbf{\Gamma}_{\text{prior}}^{1/2}$  gives

$$\begin{aligned} \mathbf{\Gamma}_{\text{post}} &= \left( \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} + \mathbf{\Gamma}_{\text{prior}}^{-1} \right)^{-1} \\ &= \mathbf{\Gamma}_{\text{prior}}^{1/2} \left( \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2} + \mathbf{I} \right)^{-1} \mathbf{\Gamma}_{\text{prior}}^{1/2}. \end{aligned} \quad (3)$$

Fast computations using this expression for  $\mathbf{\Gamma}_{\text{post}}$  require two operations: fast multiplication by  $\mathbf{\Gamma}_{\text{prior}}^{1/2}$  and fast inversion of  $\mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2} + \mathbf{I}$ . We do not address fast multiplication by  $\mathbf{\Gamma}_{\text{prior}}^{1/2}$  in this paper, since it depends on the structure of the chosen prior. Here we address the latter operation.

Let  $\lambda_i$  and  $\mathbf{v}_i$  be the eigenvalues and eigenvectors of the *prior-preconditioned Hessian of the data misfit*  $\tilde{\mathbf{H}}_{\text{misfit}} \in \mathbb{R}^{n \times n}$ , where

$$\tilde{\mathbf{H}}_{\text{misfit}} \stackrel{\text{def}}{=} \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2}.$$

Let  $\mathbf{\Lambda} \in \mathbb{R}^{n \times n}$  be the diagonal matrix of eigenvalues  $\lambda_i$  of  $\tilde{\mathbf{H}}_{\text{misfit}}$  and let  $\mathbf{V} \in \mathbb{R}^{n \times n}$  be the matrix whose columns are the eigenvectors  $\mathbf{v}_i$  of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . Then the following expression can be rewritten with  $\tilde{\mathbf{H}}_{\text{misfit}}$  replaced by its spectral decomposition:

$$\left( \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2} + \mathbf{I} \right)^{-1} = (\mathbf{V} \mathbf{\Lambda} \mathbf{V}^T + \mathbf{I})^{-1}.$$

When the eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  decay rapidly, we can extract a low-rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  by retaining only the  $r$  largest eigenvalues and corresponding eigenvectors,

$$\mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2} \approx \mathbf{V}_r \mathbf{\Lambda}_r \mathbf{V}_r^T.$$

Then we can invert using the Sherman-Morrison-Woodbury formula to obtain

$$\left( \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2} + \mathbf{I} \right)^{-1} = \mathbf{I} - \mathbf{V}_r \mathbf{D}_r \mathbf{V}_r^T + \mathcal{O} \left( \sum_{i=r+1}^n \frac{\lambda_i}{\lambda_i + 1} \right), \quad (4)$$

where  $\mathbf{\Lambda}_r \in \mathbb{R}^{r \times r}$  and  $\mathbf{V}_r \in \mathbb{R}^{n \times r}$  denote the truncated eigenvalue and eigenvector matrices, and  $\mathbf{D}_r \stackrel{\text{def}}{=} \text{diag}(\lambda_i / (\lambda_i + 1))$  where  $\mathbf{D}_r \in \mathbb{R}^{r \times r}$ . With this low-rank approximation, the final expression for approximation of the posterior covariance  $\mathbf{\Gamma}_{\text{post}}$  is therefore given by

$$\mathbf{\Gamma}_{\text{post}} \approx \mathbf{\Gamma}_{\text{prior}} - \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{V}_r \mathbf{D}_r \mathbf{V}_r^T \mathbf{\Gamma}_{\text{prior}}^{1/2}. \quad (5)$$

For many ill-posed inverse problems, the choice of  $r$  is small and independent of problem size. The  $r$  retained eigenvectors can be viewed as the modes of the parameter field that are recoverable from a combination of data and prior information. It is often the case for ill-posed inverse problems that the spectrum of the prior-preconditioned Hessian of the data misfit  $\tilde{\mathbf{H}}_{\text{misfit}}$  collapses on zero. This happens for example when the data misfit Hessian  $\mathbf{H}_{\text{misfit}}$  has rapidly decaying eigenvalues, and/or when the chosen prior is of smoothing type. In such cases, an accurate approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  can be made with small values of  $r$ . In fact, often the data and prior are informative about the low-wavenumber modes of the parameter field; this is because local features cannot be resolved from the data, and the prior is of smoothing type. Thus, refinements of the mesh on which the PDE is solved do not affect the accuracy of the low-rank approximation, and as a result, an appropriate choice of  $r$  is independent of mesh size. Note that expression (4) suggests a cutoff value for retaining eigenpairs in the low-rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$ , namely that  $\lambda \ll 1$  (in most of our examples, we use  $\lambda > 0.1$  as a cutoff criterion).

The Frobenius norm of the error in the approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  is

$$\left\| \tilde{\mathbf{H}}_{\text{misfit}} - \mathbf{V}_r \mathbf{\Lambda}_r \mathbf{V}_r^T \right\|_F = \sqrt{\sum_{j=r+1}^n \lambda_j^2}. \quad (6)$$

If the eigenvalues  $\lambda_j$  decay rapidly enough for the series  $\sum_{j=r+1}^{\infty} \lambda_j^2$  to converge, then provided the structure of the spectrum does not change as the dimension of the parameter space increases, we obtain an error bound on the low-rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  that is independent of the dimension of the parameter space (and therefore mesh size). Similarly, the error in the low-rank approximation of (4) is given by

$$\left\| \left( \tilde{\mathbf{H}}_{\text{misfit}} + \mathbf{I} \right)^{-1} - \left( \mathbf{I} - \mathbf{V}_r \mathbf{D}_r \mathbf{V}_r^T \right) \right\|_F = \sqrt{\sum_{j=r+1}^n \left( \frac{\lambda_j}{\lambda_j + 1} \right)^2} < \sqrt{\sum_{j=r+1}^n \lambda_j^2}, \quad (7)$$

and thus is also bounded independent of problem dimension, and is in fact less than the approximation error for  $\tilde{\mathbf{H}}_{\text{misfit}}$ . Finally, to relate the effect of the low-rank

approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  on the posterior covariance  $\mathbf{\Gamma}_{\text{post}}$ , one would need to know the spectral structure of the prior covariance  $\mathbf{\Gamma}_{\text{prior}}$ , as is evident from (3). In §4 and §5, we provide analytical and numerical evidence of the legitimacy of this low-rank approximation and its effect on the computed spectra of  $\tilde{\mathbf{H}}_{\text{misfit}}$  and the posterior covariance estimate for a statistical inverse problem involving a convection-diffusion PDE.

We are now in a position to state the complexity of the low-rank approximation and the resulting cost of evaluating  $\mathbf{\Gamma}_{\text{post}}$ . We certainly cannot explicitly construct  $\tilde{\mathbf{H}}_{\text{misfit}}$  and then perform a truncated SVD. Instead, we chose a matrix-free Lanczos method to find its dominant eigenvalues and corresponding eigenvectors. Lanczos requires only a matrix-vector product with  $\tilde{\mathbf{H}}_{\text{misfit}}$  at each iteration, which in turn requires (a) multiplications with  $\mathbf{\Gamma}_{\text{prior}}^{1/2}$ , and (b) actions of  $\mathbf{A}$  and  $\mathbf{A}^T$  on vectors. The latter involve forward and adjoint PDE solutions, which dominate the cost when the forward model is large-scale. Lanczos tends to require a number of iterations proportional to the dominant portion of the spectrum, which is bounded when  $\tilde{\mathbf{H}}_{\text{misfit}}$  is compact. Thus, under the conditions discussed above, the cost (in number of forward/adjoint PDE model solutions) of the low-rank approximation is small and independent of the mesh size. Once the  $r$  dominant eigenpairs have been found, storing the low-rank approximation requires  $(r+1)n$  floating point numbers, where  $n$  is the dimension of the parameter space. In §4 and §5 we illustrate the relationship between accuracy and size of  $r$  and provide evidence of mesh independence for a large-scale statistical inverse problem governed by a three-dimensional convection-diffusion PDE model.

Once the low-rank approximation has been constructed, the remaining costs to probe the posterior parameter density are linear in  $n$ . The costs are summarized in Table 1.

Calculation	cost for general $\mathbf{\Gamma}_{\text{prior}}$	cost for $\mathbf{\Gamma}_{\text{prior}} \propto \mathbf{I}$
variance: $\text{diag}(\mathbf{\Gamma}_{\text{post}})$	$(3r+1)n + r\gamma$	$(3r+1)n$
row/column: $\mathbf{\Gamma}_{\text{post}}\mathbf{e}_j$	$4rn + \gamma$	$2rn + r + 1$
multiplication: $\mathbf{\Gamma}_{\text{post}}\mathbf{v}$	$4rn + 2\gamma$	$(4r+1)n$
sample: $\exp(\mathbf{v}^T\mathbf{\Gamma}_{\text{post}}\mathbf{v})$	$(4r+2)n + 2\gamma$	$(4r+3)n$

TABLE 1

*Cost of common computations using  $\mathbf{\Gamma}_{\text{post}}$  once eigenvectors of  $\tilde{\mathbf{H}}_{\text{misfit}}$ , the prior-preconditioned Hessian of the data misfit, have been computed, where  $n$  is the dimension of  $\mathbf{x}$ ,  $r$  is the dimension of the low-rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$ , and  $\gamma$  is the cost of multiplying  $\mathbf{\Gamma}_{\text{prior}}^{1/2}$  by a vector. The last column corresponds to the case of an i.i.d. prior.*

**4. A 3D convection-diffusion inverse problem.** We now introduce a large-scale statistical inverse problem governed by three-dimensional convective-diffusive contaminant transport. The inverse problem seeks to find the initial concentration field of the contaminant and associated uncertainty, given noisy measurements of contaminant concentration at certain spatio-temporal locations or at final time, prior information on the initial contaminant, and the PDE model of the transport of the contaminant. In this section we also present analysis of a related one-dimensional convection-diffusion inverse problem to provide insight into the numerical experiments on the 3D inverse problem in §5.



**4.1. Problem description.** The forward problem is a discretization of a convection-diffusion PDE describing the evolution of the contaminant:

$$\begin{aligned} u_t - k\Delta u + \mathbf{v} \cdot \nabla u &= 0, & \Omega \times (0, T), \\ u &= u_0, & \Omega \times \{t = 0\}, \\ k\nabla u \cdot \mathbf{n} &= 0, & \partial\Omega_N \times (0, T), \\ u &= 0, & \partial\Omega_D \times (0, T), \end{aligned} \quad (8)$$

where  $u$  is the contaminant concentration,  $u_0$  is the initial contaminant concentration field,  $\mathbf{v}$  is the known wind velocity,  $k$  is the known diffusion coefficient, and  $T$  is the final time. Thus, the initial concentration  $u_0$  represents the unknown parameter  $\mathbf{x}$  and the sensor observations of concentration are the data  $\mathbf{y}_{\text{obs}}$ . Figure 1 describes the domain, boundary conditions, and velocity field for our example problem.

The code for statistical inversion extends an earlier code for the deterministic inverse problem [1, 2]. The discretization of the forward problem (8) uses a standard finite element approximation stabilized with SUPG (Streamline Upwind Petrov-Galerkin) to permit higher Peclet numbers. It uses linear tetrahedral elements and a uniform grid. Timestepping is implemented with Crank-Nicolson. Eigenvalues and eigenvectors are computed with Lanczos with periodic reorthogonalization [17]. PETSc [5, 6, 7] and SLEPc [18] are used throughout to enable parallelization.

As described in §2.2, the Bayesian formulation of a linear statistical inverse problem with Gaussian noise and prior is related to an appropriately-weighted least squares minimization problem. We choose to define our noise and prior pdf's by discretizing the infinite-dimensional functional

$$\min_{u_0} \left( \frac{\beta_{\text{noise}}}{2} \int_{\Omega} \int_0^T (u - u_{\text{obs}})^2 b(x, t) d\mathbf{x} dt + \frac{\beta_{\text{prior}}}{2} \int_{\Omega} u_0^2 d\mathbf{x} \right), \quad (9)$$

in which  $u(x)$  satisfies the convection-diffusion equation (8). We study two alternatives for the observation operator  $b(x, t)$ . For the sparse observation case, measurements are taken at distinct points  $x_j$ ,  $j = 0, \dots, m-1$ , over a time interval  $0 \leq t \leq T$ . For the final time observation case, measurements are taken over the entire domain at final time  $t = T$ . The two observation operators are defined by

$$b(x, t) = \begin{cases} \sum_j \delta(x - x_j) & \text{sparse observations,} \\ \delta(t - T) & \text{final time observations.} \end{cases}$$

Discretization of expression (9) is equivalent to the choice of Gaussians pdf's for the prior and noise,

$$\begin{aligned} \pi_{\text{prior}}(\mathbf{u}_0) &= \mathcal{N}(\bar{\mathbf{u}}_{0,\text{prior}}, \mathbf{\Gamma}_{\text{prior}}), \\ \pi_{\text{noise}}(\mathbf{e}) &= \mathcal{N}(\bar{\mathbf{e}}, \mathbf{\Gamma}_{\text{noise}}), \end{aligned}$$

with prior mean  $\bar{\mathbf{u}}_{0,\text{prior}} = \mathbf{0}$  and noise mean  $\bar{\mathbf{e}} = \mathbf{0}$ , and prior and noise covariances given by

$$\begin{aligned} \mathbf{\Gamma}_{\text{prior}} &= \frac{h^3}{\beta_{\text{prior}}} \mathbf{I}, \\ \mathbf{\Gamma}_{\text{noise}} &= \begin{cases} (h^3/\beta_{\text{noise}}) \mathbf{I} & \text{final time observations,} \\ (1/\beta_{\text{noise}}) \mathbf{I} & \text{sparse sensor observations,} \end{cases} \end{aligned}$$




FIG. 1. Time evolution of an atmospheric contaminant (deep blue) as it is transported through a model city composed of 10 buildings. The arrows represent wind velocity, and their size and color represent velocity magnitude. The mesh contains  $112 \times 112 \times 112$  hexahedra, each of which is further subdivided into 6 tetrahedral elements. The velocity field was generated by solving the steady Navier-Stokes equations with a parabolic velocity inflow (from 0 to a maximum velocity of  $v_{\max} = 1$ ), no-normal flow on the sides and top, traction-free outflow, and no-slip on the buildings and bottom. Additional parameters include density  $\rho = 1$  and viscosity  $\mu = 1$ . The maximum velocity within the domain for the chosen parameters is 2.78. Contaminant boundary conditions are zero concentration on the inflow, ground, and buildings; and zero flux on the sides, outflow and top. The images depict snapshots of the contaminant concentration as well as flow field at the following times: (a)  $T = 0$ . (b)  $T = 4$ . (c)  $T = 8$ . (d)  $T = 12$ .

where  $\beta_{\text{prior}}$  and  $\beta_{\text{noise}}$  are weighting factors and  $h$  is the mesh size.

Discretization of the infinite-dimensional least squares minimization problem (9) results in the following expression for the posterior mean (1),

$$\bar{\mathbf{u}}_{0,\text{post}} = \arg \min_{\mathbf{u}_0} \left( \frac{1}{2} (\mathbf{u} - \mathbf{u}_{\text{obs}})^T \mathbf{B}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{B} (\mathbf{u} - \mathbf{u}_{\text{obs}}) + \frac{1}{2} \mathbf{u}_0^T \mathbf{\Gamma}_{\text{prior}}^{-1} \mathbf{u}_0 \right)$$

where  $\mathbf{B}$  is a discretization of  $b(x, t)$ , and the discretized contaminant field  $\mathbf{u}$  satisfies  $\mathbf{K}\mathbf{u} = \mathbf{C}\mathbf{u}_0$ , which is the discretization of the forward convection-diffusion problem (8). Here,  $\mathbf{K}$  is the discretized convection-diffusion operator, and  $\mathbf{C}$  is the operator

that maps initial conditions to space-time. Solving this equation for  $\mathbf{u}$  given  $\mathbf{u}_0$ , we can rewrite the optimization problem in terms of  $\mathbf{u}_0$ ,

$$\bar{\mathbf{u}}_{0,\text{post}} = \arg \min_{\mathbf{u}_0} \left( \frac{1}{2} (\mathbf{A}\mathbf{u}_0 - \mathbf{B}\mathbf{u}_{\text{obs}})^T \mathbf{\Gamma}_{\text{noise}}^{-1} (\mathbf{A}\mathbf{u}_0 - \mathbf{B}\mathbf{u}_{\text{obs}}) + \frac{1}{2} \mathbf{u}_0^T \mathbf{\Gamma}_{\text{prior}}^{-1} \mathbf{u}_0 \right),$$

where  $\mathbf{A} = \mathbf{B}\mathbf{K}^{-1}\mathbf{C}$  is the parameter-to-observable map. The posterior covariance matrix of the initial concentration field is given by the inverse of the Hessian of this least squares function, which yields the expression (2) with  $\mathbf{A} = \mathbf{B}\mathbf{K}^{-1}\mathbf{C}$ . Forming the action of  $\mathbf{A}$  on a given contaminant initial condition-like vector  $\tilde{\mathbf{u}}_0$ , as required by the low-rank approximation algorithm of §3, involves lifting the initial condition to space-time ( $\mathbf{z} \leftarrow \mathbf{C}\tilde{\mathbf{u}}_0$ ), solution of the convection-diffusion problem for this source ( $\mathbf{w} \leftarrow \mathbf{K}^{-1}\mathbf{z}$ ), and extraction of the resulting contaminant concentrations at sensor locations ( $\mathbf{q} \leftarrow \mathbf{B}\mathbf{w}$ ). The action of the adjoint map  $\mathbf{A}^T$  on the resulting weighted concentrations  $\mathbf{\Gamma}_{\text{prior}}^{-1}\mathbf{q}$ , as also required in the low rank approximation, involve the transpose of the operations described above and proceed in similar fashion.

**4.2. Analysis of a model 1D convection-diffusion inverse problem.** In this section we examine the structure of an analogous 1D infinite-dimensional convection-diffusion inverse problem to assess the effectiveness of a low-rank approximation. We describe the transport of a contaminant with the following convection-diffusion initial-boundary value problem:

$$\begin{aligned} u_t - k \frac{\partial^2 u}{\partial x^2} + v \frac{\partial u}{\partial x} &= 0, & (0, L) \times (0, T), \\ u &= u_0, & (0, L) \times \{t = 0\}, \\ k \frac{\partial u}{\partial x} \Big|_{x=0} &= k \frac{\partial u}{\partial x} \Big|_{x=L}, & (0, T), \\ u|_{x=0} &= u|_{x=L}, & (0, T), \end{aligned}$$

where the velocity  $v$  is taken as constant and the domain is periodic on the interval  $(0, L)$ . We assume the observation operator takes measurements at final time  $T$  of the concentration over the entire domain  $(0, L)$ . This model problem is a one-dimensional version of (8), but with periodic boundary conditions replacing the Dirichlet and Neumann conditions. The adjoint PDE problem is equivalent to the forward problem reversed in time with negative velocity, and is given by the terminal-boundary value problem

$$\begin{aligned} -p_t - k \frac{\partial^2 p}{\partial x^2} - v \frac{\partial p}{\partial x} &= -u\delta(t - T), & (0, L) \times (0, T), \\ p &= 0, & (0, L) \times \{t = T\}, \\ k \frac{\partial p}{\partial x} \Big|_{x=0} &= k \frac{\partial p}{\partial x} \Big|_{x=L}, & (0, T), \\ p|_{x=0} &= p|_{x=L}, & (0, T). \end{aligned}$$

We wish to investigate the viability of a low-rank approximation of the data misfit Hessian matrix corresponding to (9). For the finite-dimensional inverse problem, the prior-preconditioned data misfit Hessian matrix  $\tilde{\mathbf{H}}_{\text{misfit}} = \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2}$  includes the influence of the prior and noise covariance matrices. As a surrogate for

this, we will analyze, and verify the rapid decay in, the spectrum of the infinite-dimensional operator

$$\tilde{H}_{\text{misfit}} \stackrel{\text{def}}{=} \frac{\beta_{\text{noise}}}{\beta_{\text{prior}}} A^T A$$

where  $A : u_0(x) \mapsto u(x, T)$  is the forward operator,  $A^T : u(x, T) \mapsto p(x, 0)$  is the adjoint operator, and  $\beta_{\text{noise}}/\beta_{\text{prior}}$  represents the relative strength of the prior and the noise. Provided that the (finite dimensional) prior and noise covariance matrices do not corrupt the rapid decay in the spectrum seen in  $\tilde{H}_{\text{misfit}}$ , the low-rank approximation should still be appropriate. As mentioned before, a smoothing prior often reinforces this spectral decay. The expression for the eigenvalues of  $\tilde{H}_{\text{misfit}}$  provides insight into the choice of prior and noise covariances that makes the low-rank approximation appropriate for a given statistical inverse problem. The choice of “spectrally neutral” i.i.d. prior and noise covariance matrices as in the previous subsection implies that the spectrum of  $\tilde{H}_{\text{misfit}}$  should still decay rapidly.

One can verify readily that the eigenfunctions of  $\tilde{H}_{\text{misfit}}$  are given by

$$\cos\left(\frac{2\pi jx}{L}\right) \text{ and } \sin\left(\frac{2\pi jx}{L}\right) \text{ for } j = 1, 2, \dots, \quad (10)$$

and the eigenvalues are given by

$$\frac{\beta_{\text{noise}}}{\beta_{\text{prior}}} e^{-8kT\pi^2 j^2/L^2} \text{ for } j = 1, 2, \dots \quad (11)$$

Note that the eigenvalues decay exponentially, which invites a low-rank approximation. The prior-preconditioned data misfit Hessian of the 3D inverse problem discussed in the next section is expected to exhibit similar structure, despite the effects of discretization, different boundary conditions, and complex geometry (as will be verified in the §5). Assuming the eigenvalues of  $\tilde{H}_{\text{misfit}}$  decay according to the expression (11), for a given cutoff  $\lambda > \alpha$ , the number of retained eigenvalues is therefore

$$r = 2 \left\lceil \sqrt{\frac{L^2}{8kT\pi^2}} \left| \log\left(\frac{\alpha\beta_{\text{prior}}}{\beta_{\text{noise}}}\right) \right| \right\rceil,$$

where  $0 < \alpha \leq \beta_{\text{noise}}/\beta_{\text{prior}}$ . The subsections of §5 investigate thoroughly the properties of  $\tilde{H}_{\text{misfit}}$  for the 3D numerical example and support the analysis of the infinite-dimensional  $\tilde{H}_{\text{misfit}}$  in this subsection.

We are also interested in the effect of physical parameters on the spectrum of  $\tilde{H}_{\text{misfit}}$  and the corresponding eigenfunctions. As seen in (10), the eigenfunctions of the operator  $\tilde{H}_{\text{misfit}}$  are oscillatory, where large eigenvalues are associated with smoother eigenfunctions and small eigenvalues are associated with increasingly oscillatory eigenfunctions. This is due to the effects of diffusion, which leads to greater loss of information in more oscillatory eigenfunctions (which are components of the initial concentration field), and therefore these eigenfunctions are harder to resolve from later-time observations. This behavior is explored for the 3D problem in §5.1.

Increasing the diffusion coefficient  $k$  leads to faster decay of the spectrum of  $\tilde{H}_{\text{misfit}}$ , as seen in (11). Moreover, a larger final time  $T$  also increases the rate of decay, since it allows more time for diffusion to act on components of the initial concentration. In this 1D example, there is no dependence of the spectrum of  $\tilde{H}_{\text{misfit}}$

on the velocity  $v$ , since convection of the contaminant in the adjoint operator cancels (the oppositely-signed) convection in the forward operator due to periodicity. In the 3D numerical example, effects of convection may be visible, since with a more complex velocity field and non-periodic boundary conditions, convection does not cancel out between forward and adjoint PDEs. Effects of the physical parameters  $k$  and  $T$ , as well as additional experimental parameters, on the 3D numerical spectrum are further explored in §5.2.

We next analyze the effect of the choice of truncation of the spectrum in the low-rank approximation on the resulting error in the posterior covariance matrix. To facilitate the analysis, we assume the eigenvalues of the finite-dimensional  $\tilde{\mathbf{H}}_{\text{misfit}}$  decay according to the infinite-dimensional expression (11). A bound on the squared error in the Frobenius norm (in both the approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  and the approximation of  $(\tilde{\mathbf{H}}_{\text{misfit}} + I)^{-1}$ ) in terms of  $r$  is derived by treating (6) and (7) as right Riemann sums (the factor of two is from the repeated eigenvalues). Thus,

$$\begin{aligned} 2 \sum_{j=r+1}^n \lambda_j^2 &< 2 \sum_{j=r+1}^{\infty} \lambda_j^2 < \frac{2\beta_{\text{noise}}^2}{\beta_{\text{prior}}^2} \int_r^{\infty} e^{-16kT\pi^2 x^2/L^2} dx, \\ &< \frac{2\beta_{\text{noise}}^2}{\beta_{\text{prior}}^2} \sqrt{\frac{L^2}{32kT\pi}} \operatorname{erfc} \left( r \sqrt{\frac{16kT\pi^2}{L^2}} \right). \end{aligned}$$

Alternatively, a bound in terms of the cutoff  $\alpha$  (such that retained eigenvalues satisfy  $\lambda > \alpha$ ) can be derived. Let  $R = \sqrt{\frac{L^2}{8kT\pi^2} \left| \log \left( \frac{\alpha\beta_{\text{prior}}}{\beta_{\text{noise}}} \right) \right|}$ , so that  $r = \lfloor R \rfloor$ . Then

$$\begin{aligned} 2 \sum_{j=r+1}^n \lambda_j^2 &< \frac{2\beta_{\text{noise}}^2}{\beta_{\text{prior}}^2} \int_r^{\infty} e^{-16kT\pi^2 x^2/L^2} dx, \\ &< 2 \left( \lambda_R^2 + \frac{\beta_{\text{noise}}^2}{\beta_{\text{prior}}^2} \int_R^{\infty} e^{-16kT\pi^2 x^2/L^2} dx \right), \\ &= 2 \left( \lambda_R^2 + \frac{\beta_{\text{noise}}^2}{\beta_{\text{prior}}^2} \sqrt{\frac{L^2}{32kT\pi}} \operatorname{erfc} \left( R \sqrt{\frac{16kT\pi^2}{L^2}} \right) \right), \\ &= 2 \left( \alpha^2 + \frac{\beta_{\text{noise}}^2}{\beta_{\text{prior}}^2} \sqrt{\frac{L^2}{32kT\pi}} \operatorname{erfc} \left( \sqrt{2 \left| \log \left( \frac{\alpha\beta_{\text{prior}}}{\beta_{\text{noise}}} \right) \right|} \right) \right), \\ &< 2 \left( 1 + \sqrt{\frac{L^2}{32kT\pi}} \right) \alpha^2, \end{aligned}$$

where  $0 < \alpha \leq \beta_{\text{noise}}/\beta_{\text{prior}}$ . This expression is independent of the dimension of the parameter space, and therefore the number of retained eigenvalues in the low-rank approximation need not increase to maintain accuracy as the number of parameters increases. In §5.4 we present numerical experiments that examine the effect of the error due to truncation on the posterior variance for the 3D convection-diffusion inverse problem.

**5. Numerical experiments.** We now study the behavior of the algorithm of §3 for the 3D convection-diffusion inverse problem described in §4. We examine the properties of the spectrum and eigenvectors of  $\tilde{\mathbf{H}}_{\text{misfit}}$ , and present experiments that demonstrate the effect of number of sensors, diffusion coefficient, and final time on

the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . Additionally, we investigate properties of the variance of the initial concentration and the effects of the truncation of the low-rank approximation on the computed variance. Finally, we examine the computational work associated with the low-rank approximation and demonstrate the scalability of the low-rank algorithm.

### 5.1. Properties of the spectrum of $\tilde{\mathbf{H}}_{\text{misfit}}$ and the posterior covariance.

In this section, we study the properties of eigenvalues and eigenvectors of  $\tilde{\mathbf{H}}_{\text{misfit}}$  and the resulting posterior covariance matrix.

Figure 2(a) presents the spectrum for the case of final time observations. The first 3500 eigenvalues of the spectrum of the  $\tilde{\mathbf{H}}_{\text{misfit}}$  are shown in the figure. As can be seen, the eigenvalues rapidly collapse onto zero, and for this particular problem, 1500 modes (out of 274,625) are sufficient to capture the non-trivial structure of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . The first 3500 eigenvalues of the posterior covariance matrix (based on a 3500-eigenpair low-rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$ ) are displayed in Figure 2(b). The asymptotic horizontal line of the curve reflects prior information, which in this case is a scaled standard normal pdf. The information provided by observational data is responsible for the departure from the horizontal line seen in this figure. As expressed by (5), the eigenpairs of the prior covariance matrix that are modified by observational data (i.e. the first 1500 or so) are precisely the dominant eigenpairs of  $\tilde{\mathbf{H}}_{\text{misfit}}$  (i.e. the first 1500 or so eigenvalues in Figure 2(a)). The posterior pdf thus inherits its uncertainty from the prior in those directions for which the data (filtered through the prior) render no useful information (these are eigenvectors 1500 through 274,625). The remaining directions—for which the data (filtered through the prior) do provide useful information—serve to reduce the uncertainty in the posterior.




FIG. 2. Spectra of  $\tilde{\mathbf{H}}_{\text{misfit}}$  and the posterior covariance estimate for the final time observation case with final time  $T = 6$ , diffusion coefficient  $k = 0.5$ ,  $\beta_{\text{noise}}/\beta_{\text{prior}} = 100$ , and a  $64 \times 64 \times 64$  mesh. (a) Largest 3500 (out of 274,625) eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . (b) Corresponding 3500 (out of 274,625) eigenvalues of the approximation of the posterior covariance  $\Gamma_{\text{post}}$ .

Next, we discuss the spectra for the case of sparse observations with 51 sensors and 274,625 uncertain parameters. In Figure 3 we present the spectrum and isosurfaces of selected eigenvectors of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . Figure 3(a) depicts the largest 1800 eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$ , while Figure 3(b) shows the corresponding spectrum of the posterior covariance matrix based on this low-rank approximation. Note that, like Figure 2(a), the eigenvalues decay rapidly and are negligible beyond the 500th (out of 274,625). How-

ever, the eigenvalue curve for this sparse observation case is discontinuous, exhibiting two visible jumps. The first jump occurs after the 51st eigenvalue, which is identical to the number of sensors. Indeed, each of the first 51 eigenvalues is associated with a unique sensor; this is illustrated by the 51st eigenvector shown in Figure 3(c), which depicts an isosurface in the shape of a ball surrounding the associated sensor. After the 51st eigenvalue, the associated eigenvectors change in shape by developing a large oscillation upstream of the sensor. The second jump in the spectrum occurs after the 101st eigenvalue, where the associated eigenvectors add a second oscillation upstream of the sensor. In §4.2, we noted that for a one-dimensional model problem with final-time observations, eigenvectors associated with smaller eigenvalues are increasingly oscillatory, reflecting the inability to reconstruct these modes due to diffusion-driven loss of information. This is also the case for time-dependent sparse observations within the complex geometry of the model city. However, in this case, the oscillations in the eigenvectors are not distributed evenly throughout the domain. They are focused on sensor locations, beginning with a ball surrounding a single sensor as in Figure 3(c), and becoming increasingly oscillatory and extending further upstream of the sensor, as seen in Figure 3(d)–3(f). Note that the oscillations in the eigenvectors are not uniform, likely due to the Neumann outflow boundary conditions and influence of convection on the prior-preconditioned Hessian of the data misfit.<sup>1</sup> These eigenvectors represent regions of influence for the sensors; these regions affect how the prior covariance matrix is modified in (5) to yield the posterior covariance estimate.

Similar to the final time observation case, information from observational data lowers the eigenvalues of the sparse-observation posterior covariance matrix associated with dominant eigenpairs of  $\tilde{\mathbf{H}}_{\text{misfit}}$ , thereby reducing the first 500 or so eigenvalues in Figure 3(b). However, the reduction in uncertainty (i.e. the departure from the horizontal line in the figure) is largest around the sensors themselves, as illustrated by the eigenvector depicted in Figure 3(c).

**5.2. Dependence of spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  on physical and experimental parameters.** While we argued in §3 that refining the mesh for a class of ill-posed inverse problems does not affect the dominant eigenvalues/eigenvectors of the prior-preconditioned Hessian of the data misfit, other factors do affect the rate of decay of the spectrum and therefore the number of dominant eigenvalues required for the approximation of  $\mathbf{\Gamma}_{\text{post}}$ . In this subsection, we study numerically the dependence of the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}} = \mathbf{\Gamma}_{\text{prior}}^{1/2} \mathbf{A}^T \mathbf{\Gamma}_{\text{noise}}^{-1} \mathbf{A} \mathbf{\Gamma}_{\text{prior}}^{1/2}$  on number of sensors, diffusion coefficient, and final time.

In Figure 4(a), we examine the dependence of the spectrum on the number of sensors for the sparse sensor case. As can be seen in the figure, as the number of sensors increases, the eigenvalues decay less rapidly due to the information provided by the additional sensors. Thus, more eigenpairs should be retained to maintain accuracy of the approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . Table 5.2 presents the number of eigenvalues that need to be retained for a given accuracy, and the work required to capture them. This table demonstrates that the low-rank approximation presented in §3 permits solution of the statistical inverse problem for very large numbers of parameters at very low cost. For the case of 1,442,897 parameters, fewer than 0.1% of the eigenvalues are above the cutoff of 0.1, and thus an inexpensive approximation of the posterior parameter

<sup>1</sup>Recall that in the analysis of the 1D model problem of §4.2 with periodic boundary conditions, convection in the forward and adjoint problems cancel each other, leading to the absence of convection in the data misfit Hessian.

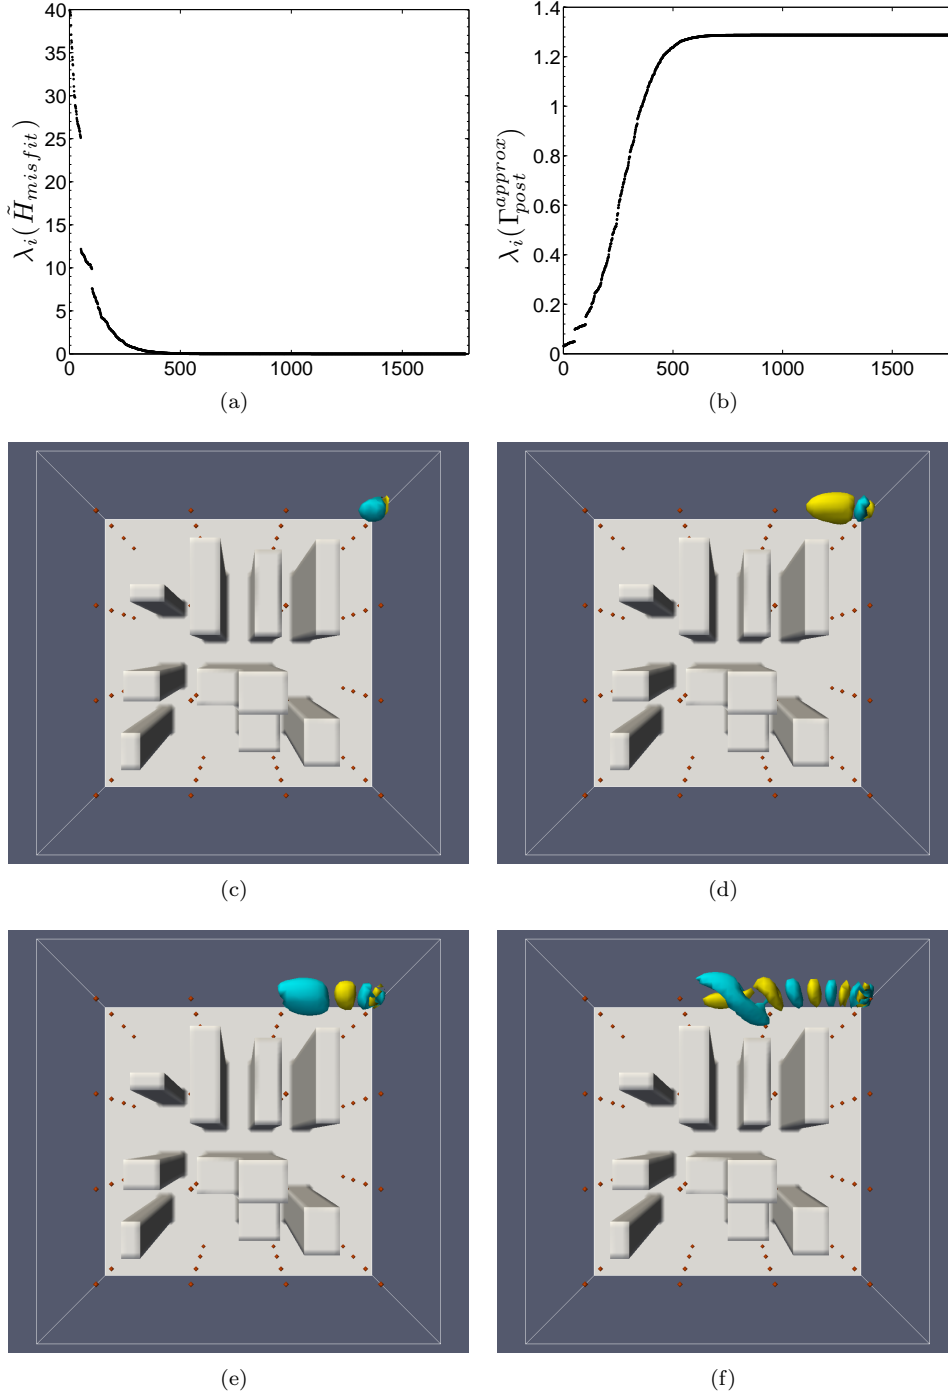


FIG. 3. (a)–(b) Spectra of  $\tilde{\mathbf{H}}_{misfit}$  and the approximated posterior covariance for a sparse sensor observation case with 51 sensors, final time  $T = 8$ , diffusion coefficient  $k = 0.05$ ,  $\beta_{noise}/\beta_{prior} = 100$ , and a  $64 \times 64 \times 64$  mesh. Largest 1800 (out of 274,625) eigenvalues of  $\tilde{\mathbf{H}}_{misfit}$ , and corresponding 1800 eigenvalues  $\Gamma_{post}$  are shown. (c)–(f) Isocontours of select eigenvectors of  $\tilde{\mathbf{H}}_{misfit}$ . Eigenvectors 51, 119, 185, and 359 are shown in (c), (d), (e), and (f), respectively. Red dots mark the locations of the sensors.



covariance can be readily made. For 27 sensors, 285 Hessian-vector products are required to capture the dominant portion of the spectrum, which as mentioned in §3 results in one forward/adjoint PDE solution per Hessian-vector product. This is dramatically fewer than the full 1.4 million forward/adjoint PDE solves needed to form the Hessian matrix, which would require years of supercomputing time. Note that while the number of dominant eigenvalues appears to increase linearly with the number of sensors, we cannot expect such a relationship in general. Here, the additional sensors are providing independent information due to the sparse placement of the sensors, while in general redundancy of sensor information may occur, permitting fewer retained eigenvalues and eigenvectors.

Figure 4(b) depicts the effect of changes in the diffusion coefficient on the spectrum, for a problem with 114 sensors. Increased diffusion results in less information captured by sensor observations, and so the spectrum decays more quickly. Thus, a lower rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  can be used to estimate  $\mathbf{\Gamma}_{\text{post}}$ .




FIG. 4. Dependence of spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  on sensor density and diffusion coefficient for the sparse observations case with final time  $T = 8$ ,  $\beta_{\text{noise}}/\beta_{\text{prior}} = 100$ , and 274,625 initial concentration parameters. (a) Dependence of spectrum on the number of sensors, for diffusion coefficient  $k = 0.05$ . (b) Dependence of spectrum on the diffusion coefficient, for a problem with 114 sensors.

Number of sensors	Number of eigenvalues	Hessian-vector products
27	209	285
51	432	585
114	871	1130

TABLE 2

Influence of number of sensors on the work (in terms of Hessian-vector products) necessary to compute a low-rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  for an eigenvalue cutoff of  $\lambda > 0.1$ . Problem has 1,442,897 initial concentration parameters, diffusion coefficient  $k = 0.05$ ,  $\beta_{\text{noise}}/\beta_{\text{prior}} = 100$ , and final time  $T = 8$ .

The final time  $T$  also affects the spectrum as shown in Figure 5. For the sparse sensor case shown in Figure 5(a), measurements are taken at each timestep until the final time  $T$ , so that an increase in  $T$  leads to an increase in information provided by the sensors, thus requiring a larger number of retained eigenpairs in the low-rank approximation. However, increasing the final time leads to diminishing returns, since as the initial contaminant field is transported out of the domain, fewer and fewer sensors provide information that can be used to infer the initial condition. Eventually,

a longer time window  $(0, T)$  will yield no new information, and the resulting curves in Figure 5(a) will tend toward each other. The diminishing returns are evident as we move from  $T = 2$  to 4 to 8 to 16 in this figure.

For the final time observation case presented in Figure 5(b), increasing  $T$  has an opposite effect on the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$ , for two reasons. First, the longer we wait to observe, the more information is lost to diffusion, (which damps modes in relation to their wave number). Second, as  $T$  increases, more of the initial contaminant field is transported out of the range of the sensors, so that the amount of information provided by the sensors at final time  $T$  diminishes. These effects are evident in Figure 5(b). Note that this final time observation case most closely resembles the analytic results presented for the continuous 1D model problem in §4.2, although in that case the periodic boundary conditions prevented information loss due to convection.

Since a change in  $T$  affects all of the sensor data in the final time observation case, while it affects only some of the data in the sparse observation case, the influence of  $T$  on the spectrum is more pronounced in the former case, as can be seen by comparing Figures 5(a) and 5(b).




FIG. 5. Effects of changing the final time  $T$  on the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  for the two sensor observation scenarios, for a problem with  $k = 0.05$ ,  $\beta_{\text{noise}}/\beta_{\text{prior}} = 100$ , and 274,625 parameters. (a) Sparse sensor observation case (with 27 sensors). (b) Final time observation case.

**5.3. Interpretation of the posterior variance.** The approximation of the posterior covariance is given by (5). For the present example, the specific choice of  $\mathbf{\Gamma}_{\text{prior}}$  and  $\mathbf{\Gamma}_{\text{noise}}$  are given in §4.1 as scaled identities, leading to

$$\mathbf{\Gamma}_{\text{post}} \approx \frac{h^3}{\beta_{\text{prior}}} (\mathbf{I} - \mathbf{V}_r \mathbf{D}_r \mathbf{V}_r^T).$$

Recall that  $\mathbf{V}_r$  and  $\mathbf{D}_r$  represent the dominant eigenpairs of the prior-preconditioned Hessian matrix of the data misfit,  $\tilde{\mathbf{H}}_{\text{misfit}}$ .

In Figure 6, we plot the approximate posterior variance of the initial concentration field, which is given by the diagonal of  $\mathbf{\Gamma}_{\text{post}}$ , for the case of 51 sensors distributed uniformly (avoiding the building interiors) throughout the model city. Figure 6(a) presents a cross-section of the variance field, taken horizontally through the top row of sensors. Recall that the prior estimate of the variance is  $h^3/\beta_{\text{prior}}$ , and the posterior estimate modifies this value according to the expression (5). The bright red area is the

region of the domain where the prior dominates. Information gained from observations does not reduce the uncertainty in the contaminant concentration in these regions. The influence of the observations on the posterior variance is concentrated around the sensors, as expected. The lowest uncertainty (dark blue) is exactly at the sensors, with increasing uncertainty moving away from the sensor locations. A three-dimensional view of the variance is given in Figure 6(b), which shows isocontours of the variance in light blue in the context of the surrounding buildings, superposed on streamlines of the velocity field. The isocontours capture the low uncertainty region surrounding each sensor. Note that these low variance regions vary in shape. In the leeward regions of buildings, where diffusion dominates, the low variance regions are more isotropic. Where convection dominates, the low variance regions associated with each sensor extend anisotropically upstream, following the streamlines of the velocity field. Contaminant that is initially directly upstream of a sensor is convected along streamlines and measured by downstream sensors, thereby reducing uncertainty of the reconstructed initial concentration in these upstream regions. On the other hand, contaminants downstream of a sensor tend to experience greater diffusion as they are convected a further distance before encountering a downstream sensor. Thus, the stronger the convection, the longer the low variance tails. In low-velocity regions where contaminants spread mainly through diffusion, however, the variance reduces in a more isotropic fashion. The alignment of the anisotropy of the variance with streamlines is evident in Figure 6.

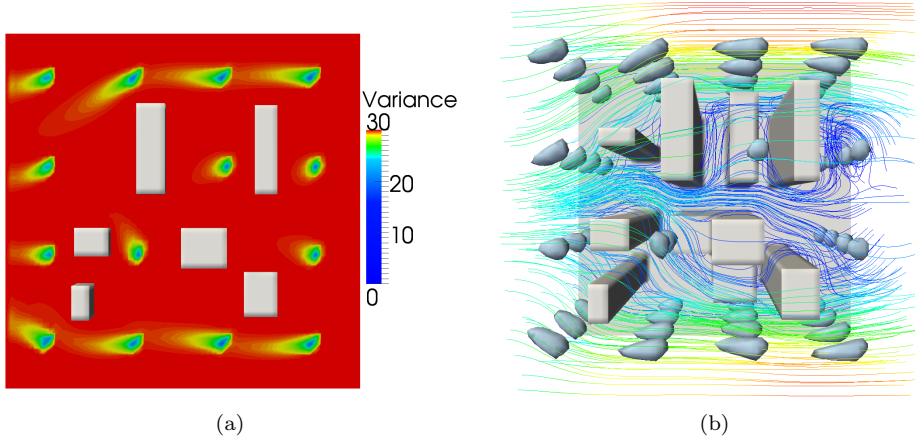


FIG. 6. The variance in the initial concentration of the contaminant, for the sparse observation case of 51 sensors measuring to final time  $T = 8$  with diffusion coefficient  $k = 0.03$  for a  $64 \times 64 \times 64$  mesh. The covariance approximation is based on an eigenvalue cutoff of  $\lambda > 0.1$ , and required 513 eigenvalues out of 274,625, computed at a cost of 650 forward/adjoint convection-diffusion PDE solutions. We chose  $\beta_{\text{prior}} = h^3/30$  and  $\beta_{\text{noise}} = 100\beta_{\text{prior}}$ . (a) Cross-section of the variance field through the top row of sensors, with arrows representing the velocity field. (b) Top view, isocontours of the variance field superposed on streamlines of the velocity field that are colored by magnitude.

**5.4. Effect of low-rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$  on accuracy of variance approximation.** In this section, we study the effect of truncation of the spectrum, and resulting low-rank approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$ , on the solution of the statistical inverse problem, specifically the evaluation of the main diagonal of the posterior covariance matrix, i.e. the variances of the initial concentration parameters. Figure 7 presents the relative error in the approximated variance as a function of the number of

eigenvectors retained. As can be seen in the figure, the relative error decreases rapidly with the number of retained eigenvectors of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . For both the sparse observation case and the final time observation case, inclusion of additional eigenpairs in the approximation of  $\Gamma_{\text{post}}$  has diminishing returns with respect to increased accuracy.




FIG. 7. Dependence of relative error in the approximated variance on the number of eigenvalues retained in the approximation. Since the dimension of  $\tilde{\mathbf{H}}_{\text{misfit}}$  is 274,625, computing the exact variance at each grid point,  $\text{Var}_i^{\text{exact}}$ , is prohibitive; instead, the “exact” solution is defined using an aggressive eigenvalue cutoff of  $\lambda > 10^{-5}$ . The relative error is defined by  $\sum_i |\text{Var}_i^{\text{exact}} - \text{Var}_i^{\text{approx}}| / \sum_i |\text{Var}_i^{\text{exact}}|$ . (a) Relative error in the variance for a sparse observation case with 27 sensors, final time  $T = 16$ ,  $\beta_{\text{noise}}/\beta_{\text{prior}} = 100$ , and diffusion coefficient  $k = 0.05$ . The “exact” variance retains 462 eigenpairs in the approximation of  $\tilde{\mathbf{H}}_{\text{misfit}}$ . (b) Relative error in the variance for a final time observation case with final time  $T = 6$  and diffusion coefficient  $k = 0.05$ . The “exact” variance retains 3,240 eigenpairs.

These diminishing returns are also illustrated in Figure 8, which displays the effects of truncation on the approximate variance field (on a horizontal plane) for a sparse observation case with 27 sensors. This is the same model city as the one used to generate Figure 7(a). The sequence of three images in the figure show the improvement in the approximation of the variance field as the number of retained eigenvectors is increased corresponding to a eigenvalue cutoff of 5, 1, and 0.1. The differences are most clearly visible in the increased length of the low variance tails (yellow/green) upstream of the sensors. These are the regions for which the observations still provide some information (when filtered through the prior), before yielding to the red areas, which are determined only by the prior. The lowest variance areas (in blue, at the sensor locations) extend further for the smallest cutoff value. While the 59 additional eigenvectors used in the variance approximation of Figure 8(b) compared to Figure 8(a) provide visible improvement, the additional 92 eigenvectors needed to generate the approximate variance field in Figure 8(c) provide negligible returns in accuracy.

**5.5. Dependence of the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  on mesh size and resulting scalability of the low-rank algorithm.** This section addresses the two properties required for scalability of the algorithm for low-rank-based approximation of the posterior covariance that was presented in §3. First, the number of eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  that materially influence the posterior covariance should be independent of mesh size. Second, the work to compute those eigenpairs, measured in number of forward/adjoint solves of the 3D convection-diffusion PDE problems, should be independent of mesh




FIG. 8. The variance in the initial concentration of the contaminant, based on three approximations to  $\Gamma_{\text{post}}$  from different truncation cutoff values for a sparse observation case with 27 sensors, final time  $T = 16$ , diffusion coefficient  $k = 0.05$ , and 274,625 parameters. (a) Cutoff  $\lambda > 5$ : 66 eigenvalues. (b) Cutoff  $\lambda > 1$ : 125 eigenvalues. (c) Cutoff  $\lambda > 0.1$ : 217 eigenvalues.

size. Figure 9(a) presents the dependence of the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  on the mesh size, which is the same as the size of the parameter space. The figure shows the eigenvalue behavior for five successively finer mesh sizes, corresponding to a number of initial condition parameters ranging from 117,649 to 1,442,897. As can be seen in the figure, increasing the size of the parameter space does not affect the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  substantially, and therefore does not change appreciably the number of eigenvalues required to approximate  $\Gamma_{\text{post}}$  up to a chosen eigenvalue cutoff. Indeed, a log-linear plot as in Figure 9(b) is needed to see the differences between the five spectra (note that the figure shows only the largest 0.07% to 0.9% of eigenvalues, depending on mesh size). In this figure, we see a small change in the spectrum as the mesh is refined for the smallest eigenvalues displayed in the plot, which as discussed in §5.1 are associated with more oscillatory eigenvectors. Eigenvectors associated with larger eigenvalues are smoother, and thus their corresponding eigenvalues are not affected by mesh refinement. On the other hand, the more oscillatory eigenvectors associated with smaller eigenvalues become better resolved with mesh refinement. The mesh should be fine enough to capture the eigenpairs that have an appreciable influence on the posterior covariance (and therefore should be retained in the low-rank approximation). The eigenvalue cutoff criterion provided in §3, namely that  $\lambda \ll 1$ , precisely achieves this goal. Figure 10(a), which displays the number of eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  above a cutoff of  $\lambda > 0.1$  as a function of number of parameters (and hence mesh size), demonstrates that the number of eigenvalues required to capture the dominant portion of the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  is independent of problem size, once the mesh is sufficiently fine to resolve the important (as identified from the data, filtered through the prior) eigenvectors.

The other property needed to insure scalability of our algorithm for low-rank estimation of the posterior covariance is that the work required to capture the dominant eigenvalues also does not depend on mesh size. Towards this end, we choose to extract the dominant eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  using the Lanczos algorithm option in SLEPc. The main computational kernel in the Lanczos algorithm is to form a product of the matrix  $\tilde{\mathbf{H}}_{\text{misfit}}$  with a vector. As discussed in §3,  $\tilde{\mathbf{H}}_{\text{misfit}}$  is never explicitly formed; instead, each matrix-vector product is dominated by one forward and one adjoint solve of the convection-diffusion PDE. The number of matrix-vector products is thus a good measure of the cost of the low-rank approximation: solution of the forward and adjoint PDEs overwhelms all other components, both in the eigenvalue computation and in the subsequent linear algebra to estimate the posterior covariance matrix.




FIG. 9. Dependence of the spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  on mesh refinement for a final time observation case with  $T = 6$ ,  $\beta_{\text{noise}}/\beta_{\text{prior}} = 100$ , and diffusion coefficient  $k = 0.05$ . (a) Eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  such that  $\lambda > 0.1$  for different mesh resolution. (b) Log-linear plot of eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  such that  $\lambda > 0.1$  for different mesh resolution.

These components are at most linear in the number of parameters (see Table 1 for costs of the latter). Figure 10(b) depicts the number of matrix-vector products (and therefore forward/adjoint PDE solves) to compute the retained eigenvalues shown in Figure 10(a). The parameters again range in number from 117,649 to 1,442,897. As can be seen, the number of required matrix-vector products is again largely independent of parameter dimension. In summary, Figure 10(b) illustrates that the posterior covariance matrix, and hence solution to the convection-diffusion statistical inverse problem, can be estimated accurately at a cost (measured in PDE solves) that is independent of problem size.




FIG. 10. Work required to estimate dominant spectrum of  $\tilde{\mathbf{H}}_{\text{misfit}}$  as a function of mesh size/number of parameters, for a final time observation case with  $T = 6$ ,  $\beta_{\text{noise}}/\beta_{\text{prior}} = 100$ , and diffusion coefficient  $k = 0.05$ . (a) Number of eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  retained, such that  $\lambda > 0.1$ . (b) Number of matrix-vector products (and hence forward/adjoint PDE solves) required to compute eigenvalues of  $\tilde{\mathbf{H}}_{\text{misfit}}$  such that  $\lambda > 0.1$ .

**6. Conclusions.** We have presented a scalable method for quantifying uncertainties in linear statistical inverse problems that are governed by expensive forward models (such as PDEs) and high-dimensional parameter spaces (as arise from discretization of infinite-dimensional fields). We have considered the Bayesian framework with Gaussian noise and prior, for which the solution to the inverse problem is also Gaussian and thus characterized by the mean and covariance matrix of the posterior probability density. For such problems, direct computation of the posterior covariance matrix is prohibitive. For many ill-posed inverse problems, however, the Hessian of the data misfit term has compact structure, which admits a low-rank approximation with minimal error. For such problems, we present a fast method for computation of an approximation to the covariance of the posterior that exploits the low-rank structure of the prior-preconditioned Hessian of the data misfit  $\tilde{\mathbf{H}}_{\text{misfit}}$ . This permits quantification of uncertainty in large-scale ill-posed inverse problems at a small multiple (independent of the problem dimension) of the cost of solving the forward problem. Analysis of a 1D model convection-diffusion inverse problem and numerical experiments with a large-scale 3D convection-diffusion inverse problem with up to 1.5 million parameters illustrate the scalability and accuracy of the method.

The work presented in this article addresses the class of inverse problems with linear parameter-to-observable map and Gaussian prior and additive noise probability densities. However, our method is more generally applicable to any posterior density that is well-approximated by a Gaussian. In such cases, the Hessian of interest is the Hessian evaluated at the MAP point, and the approach presented here applies directly. For non-Gaussian posteriors, the (inverse of the) Hessian of the negative log posterior (evaluated at the MAP point) is no longer equivalent to the posterior covariance matrix. For high dimensional parameter spaces, one must then resort to sampling the posterior for example by Markov chain Monte Carlo methods, either directly or by first approximating the posterior by a Gaussian process surrogate. In these cases, the method presented here may offer opportunities to substantially improve the speed and effectiveness of MCMC as well as Gaussian process construction by incorporating an approximation of the (local) prior-preconditioned Hessian of the data misfit. Results demonstrating these ideas are forthcoming [12, 13].

**Acknowledgments.** We thank George Biros, James Martin, Youssef Marzouk, and Karen Willcox for fruitful discussions related to this work.

#### REFERENCES

- [1] V. AKÇELİK, G. BIROS, A. DRAGENESCU, J. HILL, O. GHATTAS, AND B. VAN BLOEMEN WAANDERS, *Dynamic data-driven inversion for terascale simulations: Real-time identification of airborne contaminants*, in Proceedings of SC05, Seattle, Nov. 2005, IEEE/ACM.
- [2] VOLKAN AKÇELİK, GEORGE BIROS, OMAR GHATTAS, KEVIN R. LONG, AND BART VAN BLOEMEN WAANDERS, *A variational finite element method for source inversion for convective-diffusive transport*, Finite Elements in Analysis and Design, 39 (2003), pp. 683–705.
- [3] V. AKÇELİK, G. BIROS, O. GHATTAS, J. HILL, D. KEYES, AND B. VAN BLOEMAN WAANDERS, *Parallel PDE constrained optimization*, in Parallel Processing for Scientific Computing, M. Heroux, P. Raghaven, and H. Simon, eds., SIAM, 2006.
- [4] G. BACKUS AND F. GILBERT, *The resolving power of gross earth data*, Geophysical Journal of the Royal Astronomical Society, 16 (1968), pp. 169–205.
- [5] S. BALAY, K. BUSCHELMAN, V. ELJKHOUT, W. D. GROPP, D. KAUSHIK, M. G. KNEPLEY, L. C. MCINNEN, B. F. SMITH, AND H. ZHANG, *PETSc Users Manual*, Tech. Report ANL-95/11 - Revision 3.0.0, Argonne National Laboratory, 2008.
- [6] S. BALAY, K. BUSCHELMAN, W. D. GROPP, D. KAUSHIK, M. G. KNEPLEY, L. C. MCINNEN, B. F. SMITH, AND H. ZHANG, *PETSc Web page*, 2009. <http://www.mcs.anl.gov/petsc>.

- [7] SATISH BALAY, WILLIAM D. GROPP, LOIS CURFMAN MCINNES, AND BARRY F. SMITH, *Efficient management of parallelism in object oriented numerical software libraries*, in Modern Software Tools in Scientific Computing, E. Arge, A. M. Bruaset, and H. P. Langtangen, eds., Birkhäuser Press, 1997, pp. 163–202.
- [8] J. G. BERRYMAN, *Analysis of approximate inverses in tomography I. Resolution analysis of common inverses*, Optimization and Engineering, 1 (2001), pp. 87–115.
- [9] ———, *Analysis of approximate inverses in tomography II. Iterative inverses*, Optimization and Engineering, 1 (2001), pp. 437–473.
- [10] L. T. BIEGLER, O. GHATTAS, M. HEINKENSCHLOSS, AND B. VAN BLOEMEN WAANDERS, eds., *Large-Scale PDE-Constrained Optimization*, Lecture Notes in Computational Science and Engineering, Vol. 30, Springer-Verlag, Heidelberg, 2003.
- [11] L. BOSCHI, *Measures of resolution in global body wave tomography*, Geophysical Research Letters, 30 (2003).
- [12] TAN BUI-THANH, OMAR GHATTAS, DAVID HIGDON, ZHENG DONG LU, AND LUCAS WILCOX, *Hessian-informed Gaussian process response surface methods for probability density approximation*. Manuscript, 2009.
- [13] CARSTEN BURSTEDDE, OMAR GHATTAS, JAMES MARTIN, AND LUCAS C. WILCOX, *Uncertainty quantification in inverse problems with stochastic Newton MCMC*. Manuscript, 2009.
- [14] S. FOMEL, J. BERRYMAN, R. CLAPP, AND M.F. PRUCHA, *Iterative resolution estimation in least squares Kirchhoff migration*, Geophysical Prospecting, 50 (2002), pp. 577–588.
- [15] W. P. GOUVEIA AND J. A. SCALES, *Resolution of seismic waveform inversion: Bayes versus Occam*, Inverse Problems, 13 (1997), pp. 323–349.
- [16] ———, *Bayesian seismic waveform inversion: Parameter estimation and uncertainty analysis*, Journal of Geophysical Research, 103 (1998), pp. 2759–2779.
- [17] V. HERNANDEZ, J. E. ROMAN, A. TOMAS, AND V. VIDAL, *Lanczos methods in SLEPc*, Tech. Report STR-5, Universidad Politécnica de Valencia, 2006. Available at <http://www.grycap.upv.es/slepc>.
- [18] V. HERNANDEZ, J. E. ROMAN, AND V. VIDAL, *SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems*, ACM Transactions on Mathematical Software, 31 (2005), pp. 351–362.
- [19] JARI KAIPIO AND ERKKI SOMERSALO, *Statistical and Computational Inverse Problems*, vol. 160 of Applied Mathematical Sciences, Springer-Verlag, New York, 2005.
- [20] O. MARQUES, T. DRUMMOND, AND D. VASCO, *A computational strategy for the solution of large linear inverse problems in geophysics*, Parallel and Distributed Processing Symposium, International, (2003), p. 16b.
- [21] S. MINKOFF, *A computationally feasible approximate resolution matrix for seismic inverse problems*, Geophysical Journal International, 126 (1996), pp. 345–359.
- [22] G. NOLET, R. MONTELLI, AND J. VIRIEUX, *Explicit, approximate expressions for the resolution and a-posteriori covariance of massive tomographic systems*, Geophysical Journal International, 138 (1999), pp. 36–44.
- [23] ALBERT TARANTOLA, *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, Philadelphia, PA, 2005.
- [24] D. W. VASCO, JR. J. E. PETERSON, AND E. L. MAJER, *Resolving seismic anisotropy: Sparse matrix methods for geophysical inverse problems*, Geophysics, 63 (1998), pp. 970–983.
- [25] D. W. VASCO, R. J. PULLIAM, AND L. R. JOHNSON, *Formal inversion of ISC arrival times for mantle P-velocity structure*, Geophysical Journal International, 113 (1993), pp. 586–606.
- [26] C. R. VOGEL, *Computational methods for inverse problems*, Frontiers in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2002.
- [27] Z. S. YAO, R. G. ROBERTS, AND A. TRYGGVASON, *Calculating resolution and covariance matrices for seismic tomography with the LSQR method*, Geophysical Journal International, 138 (1999), pp. 886–894.
- [28] H. ZHANG AND C. H. THURBER, *Estimating the model resolution matrix for large seismic tomography problems based on Lanczos bidiagonalization with partial reorthogonalization*, Geophysical Journal International, 170 (2007), pp. 337–345.
- [29] J. ZHANG AND G. A. MCMECHAN, *Estimation of resolution and covariance for large matrix inversions*, Geophysical Journal International, 121 (1995), pp. 409–426.